

System Software and Architecture

Thanks to Tara McQueen, Paidá Munhutu, and
Remik Ziemiński for figures and assistance

Amy Langenhorst
Summer School 2012

Outline: Software & Architecture

- Environment Modules
- MOAB
- gcp
- Gaea Overview
- Pan Overview

Environment Modules

- Help users manage their shell environment
- Allow groups of related environment-variable settings to be made or removed dynamically
- Each module contains the information needed to configure the shell for an application
- Modules can be loaded and unloaded dynamically and atomically
- Modules are useful for managing multiple versions of software

Key Module Commands

- module avail
 - List all available packages
- module list
 - List currently loaded packages
- module load *package-name*
 - Insert a package into your environment
- module unload *package-name*
 - Remove a package from your environment

MOAB

- Job Scheduler
- Actually a **meta-scheduler** or **workflow manager**
 - Moab sits above a set of resource managers and allows the workflow to move between them
 - Allows compute at Gaea, postproc at GFDL
 - You will use MOAB commands for scheduling
 - Don't interact directly with the resource managers (Torque)


Key MOAB Commands

- `msub script-name`
 - Submit a job script to be scheduled
- `showq -n -v -u user`
 - List your jobs in the system
- `checkjob -v job-id`
 - Show details about a specific job
- `canceljob job-id`
 - Cancel a specific job

gcp

- “general copy” utility
- Wrapper tool which chooses optimal transfer protocol and settings between GFDL’s sites and filesystems
- Use “smart site” prefixes *gaea:* or *gfdl:* for cross-site transfers. Examples from gaea:
 - `gcp /lustre/fs/scratch/$USER/file gfdl:/archive/$USER/`
 - `gcp gfdl:/archive/$USER/file /lustre/ltfs/scratch/$USER/`
- Please use gcp instead of other tools to move data between filesystems

Gaea Architecture



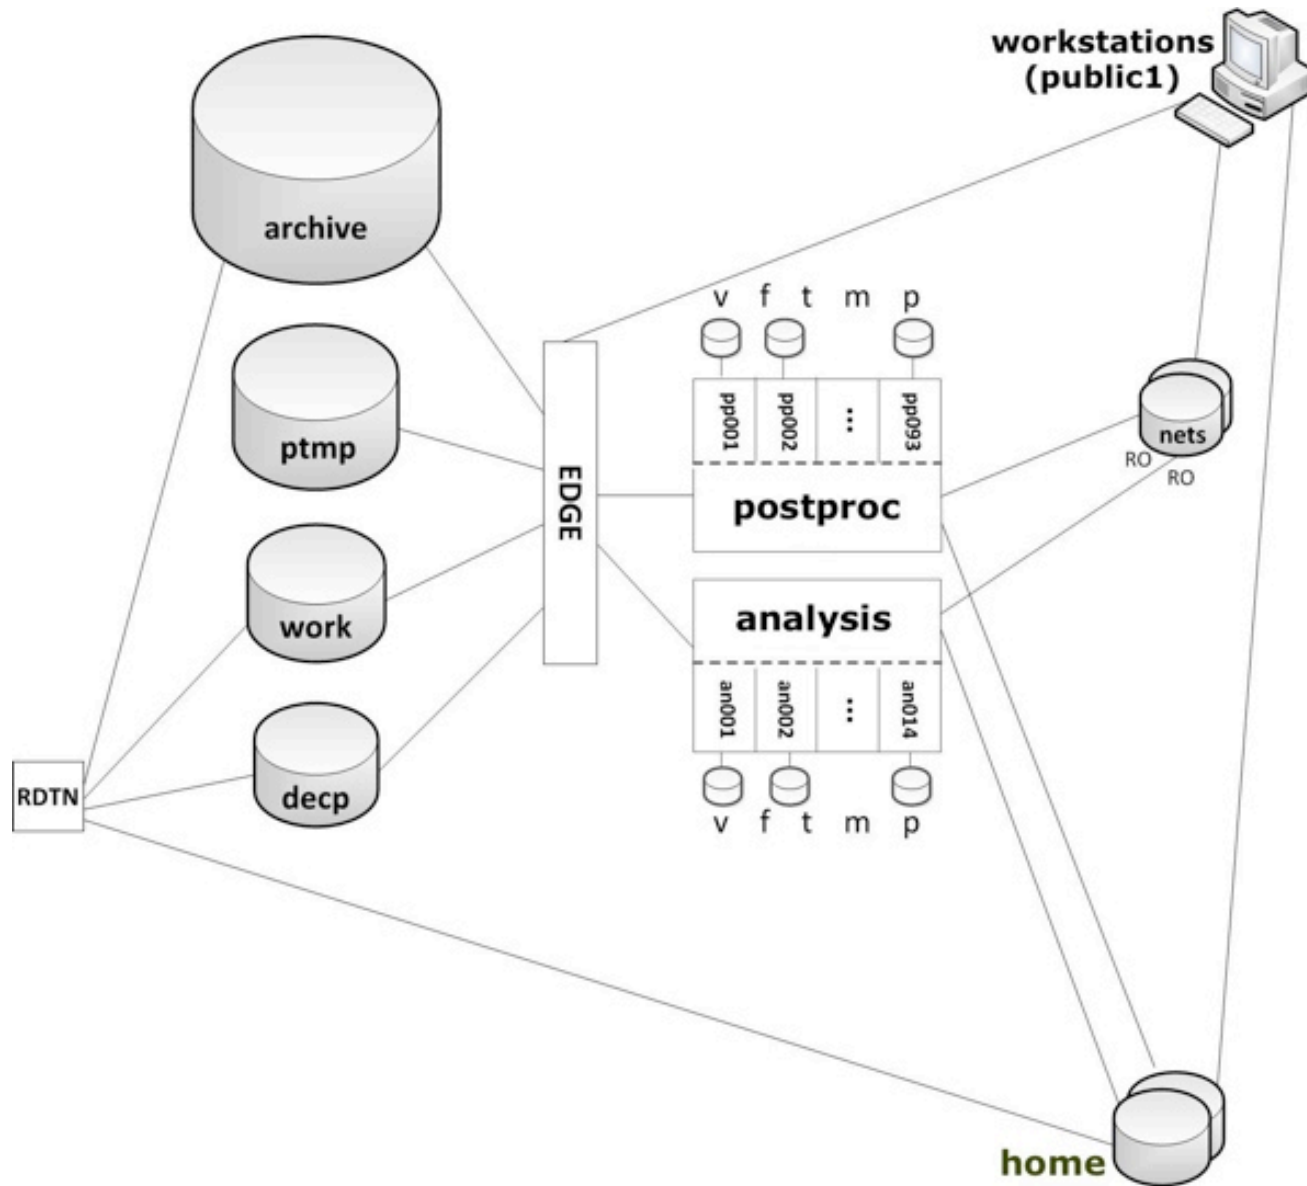
Gaea Do's

- Use gcp for transfers
- Compile on login nodes
- Put input data, source code and commonly used files on LTFS
- Put transient data on FS

Gaea Don'ts

- Don't run “module purge”
- Don't run unalias *
- Don't perform deep, large scale use of “find” on lustre filesystems
- Don't do operations like ls -R
- Don't use cp
- Don't do transfers or intensive I/O on batch/compute nodes
- Don't compile on batch/compute nodes
- Don't use FS as permanent storage

PAN Architecture



PAN Do's

- Use “module load analysis_dujour”
 - Puts latest analysis packages in your environment
- Move files using gcp
- Check whether your files in /archive are on tape or disk with “dmls -l *file*”
- If you need to bring many /archive files to disk, use “dmget *filepattern*”

PAN Don'ts

- Don't work directly in archive, it is for storage only
 - Instead, from the analysis nodes, gcp your file to \$TMP and work there
 - Be sure to gcp data you want to save back to /archive before you log off!
- Don't use cp

Questions?

- Search the GFDL Wiki for more on these topics



Image by
Remik
Ziemlinski